

REMARKS

Procedural History and Summary of Office Action

This application was finally rejected on June 19, 2006. This reply is being filed with a Request for Continued Examination (37 C.F.R. § 1.114). Accordingly, the finality of the rejection should be withdrawn and prosecution should continue.

Claims 61-80 were pending in the above-identified patent application.

The Examiner rejected claims 61-80 under 35 U.S.C. § 103(a) as being obvious from Farber et al. U.S. Patent 6,185,598 in view of Feigenbaum U.S. Patent 6,339,785.

Summary of Applicants' Reply

Applicants have amended claims 61-63, 65, 75, 76 and 80 in order to more particularly define the invention. The Examiner's rejection is respectfully traversed.

Summary of Personal Interview

On August 14, 2006, the undersigned and Mr. Dar Efroni, CTO of applicants' assignee, conducted a personal interview with the Examiner and his colleague Examiner Jeffrey C. Pwu, at the Patent and Trademark Office. The undersigned, on behalf of Mr. Efroni and himself, would like to thank the Examiners for the courtesies extended during the personal interview. The substance of the interview will be discussed below as appropriate.

Applicants' Reply to the Prior Art Rejection

Claims 61-80 were finally rejected under 35 U.S.C. § 103(a) as being obvious from Farber in view of Feigenbaum. This rejection is respectfully traversed.

Applicants' invention, as defined by the pending claims, as amended, is a caching ("acceleration") method or server for a peer-to-peer ("client-to-client") network. Although applicants' specification equates the terms "peer-to-peer" and "client-to-client" (see specification, page 2, line 4), applicants have amended the claims to replace each occurrence of "client-to-client" with "peer-to-peer" in order to clarify the claims by using the more common term,* and thereby emphasizing the distinction over the references.

In accordance with the invention as claimed, all data traffic to and from a first client in the network is monitored for cacheable queries to multiple other clients and responses thereto, which come from multiple other clients. When such queries or responses are identified, they are intercepted by the acceleration server and the responses aggregated before being transmitted to the first client (and potentially later to other clients). The acceleration server and the method carried out thereby, and particularly the interception of queries and responses by the acceleration server, and the transmission of responses by the acceleration server, are transparent to the clients, meaning that the clients continue to believe that they are communicating directly with each other.

Neither Farber nor Feigenbaum, whether taken separately or together, shows or suggests the claimed invention.

Farber shows what is essentially a system for mirroring busy servers, except that instead of full mirror

* Applicants respectfully submit that because it merely clarifies what was already in the claim by replacing one term with an equivalent term, this amendment is not a substantial amendment, or a narrowing amendment, under the doctrine of Festo Corp. v. Shoketsu Kinzoku Kogyo Kabushiki Co., 535 U.S. 722, 62 USPQ2d 1705 (2002), and related cases.

servers, "repeaters" that "sparsely" mirror the server are provided. Requests addressed to the server are "reflected" by "reflectors" from the server to one of the repeaters, depending on load conditions. If the repeater to which the request is directed does not have the requested data (because it is only a sparse mirror), it will obtain it from the server, or possibly from another repeater, before delivering it to the requester (thus over time the repeater will become less sparse and more like a mirror).

Feigenbaum shows a client-server network for downloading files, in which a client requests particular data, and different portions of that data are delivered by different servers. There is no caching or acceleration disclosed or suggested by Feigenbaum, nor does Feigenbaum show a peer-to-peer network.* Although Feigenbaum does not use the term "mirror," the multiple servers in Feigenbaum may be considered to be mirrors.

Thus, both Farber and Feigenbaum show systems where files may be downloaded from mirrored servers, and they do not show peer-to-peer networks.

As previously explained by applicants, the need to monitor traffic in both directions distinguishes peer-to-peer network traffic from a simple download request to a server from a client, because the response to a peer-to-peer request can come from any number of other peers (whether considered servers or clients), and also because the requesting client can itself be a server of some other request from a different client. Moreover, as explained during the interview, in peer-

* Applicants erred in the previous Reply to Office Action in characterizing Feigenbaum as a peer-to-peer network. Feigenbaum does not show two-way or multi-way downloading, but only downloading from redundant servers designed to have data downloaded from them. Thus, applicants do not admit that Feigenbaum shows a peer-to-peer network.

to-peer communications it is important that the participants -- the clients of the claims, including both the "first client" and the "at least two other clients" -- continue to believe at all times that they are communicating with each other, because peer-to-peer protocols give "credit" to clients that provide response portions, elevating those clients' standing in terms of receiving responses to their own queries. This means that when the acceleration server delivers to the first client data that it had cached from one of the other clients, the first client has to believe that the data came from that other client. The clients cannot know that their link was intercepted or that the transmission of data to the first client was from the acceleration server.

Applicants have amended the claims to define this invisibility or transparency. In particular, in method claims 61 and 65, the interception (which in claim 65 includes redirection) and transmission are defined as being performed transparently, and machine-readable storage medium claim 80, which is based on claim 65, includes a corresponding definition. Similarly, in apparatus claims 75 and 76, the acceleration server and the redirecting device are defined as being transparent to the clients.

These amendments are fully supported by the application as filed. The characterization of the acceleration server as being transparent is set forth explicitly in the specification as filed in multiple locations, including at least page 6, lines 10-11, page 16, lines 2-3, and original claims 26 and 55 as filed (which form part of the original disclosure). That disclosure is further enabled by the disclosure of the use of a redirecting device (including the disclosure of using a Level 4 switch) for interception at page 3, lines 6-10, and page 11, lines 4-6, and in original claims 3, 4 and 32 as filed. In addition, the disclosure at page 13, line 7, through page 14, line 4, of

returning cached data only when the client that is the source of the cached data is actually currently on-line, and in some embodiments only if that client actually still has the cached data (even though the data is coming from the cache), supports the transparent operation of the invention in a manner that the requesting client ("first client") believes that it is communicating with the other client.

There is no disclosure in either Feigenbaum or Farber of the claimed transparency.

In Feigenbaum, there is no intermediate device that acts as the claimed monitoring and intercepting mechanism (or that performs the claimed monitoring or intercepting functions) that can even be made transparent (or not transparent). Moreover, somehow the requesting client determines that it must send its request to several different servers (Feigenbaum, column 3, lines 19-32), so nothing is transparent in Feigenbaum.

In Farber, the reflector, after choosing a repeater to service a request in place of the origin server, does not tell the repeater to service the request. Rather, it sends a redirect message back to the requesting client with the identity of the repeater (Farber, column 3, lines 10-12, and column 8, lines 50-51). The client then contacts the repeater directly (column 13, lines 13-16, and column 10, lines 14-19). In this case, the existence of the reflector and the repeater is clear to the client, and the interception of the request is explicit, both are which are the exact opposite of being transparent.*

* Applicants note that at column 5, line 15, Farber uses the term "transparently" to describe the reflection of a request to a repeater. In view of the remaining description in Farber, applicants do not see how the process of Farber can be considered transparent in any way. What is clear, however, is that whatever Farber means by "transparently," it does not (continued...)

In addition, the references fail for other reasons to meet the limitations of applicants' claims.

As stated above, in Feigenbaum, there is no intermediate device that acts as the claimed monitoring and intercepting mechanism, or performs the claimed monitoring or intercepting functions. Instead, somehow the requesting client determines that it must send its request to several different servers.

There is also no such intermediate device in Farber. The Examiner has taken the position that the repeater of Farber has to monitor traffic in both directions because it has to respond to queries, and also has to capture data that are included in responses from the main server in cases where the repeater does not have the data (so that the repeater can store it for future use). However, that is not enough for applicants' claims to read on Farber.

Taking claim 75 as representative of applicants' independent claims, the claim defines, in the context of a transaction, and in addition to the transparency discussed above, a first client that implements a query to "other clients" and those "other clients" provide response portions to the first client that answer the first client's query. An acceleration server includes a monitoring mechanism that monitors all communications to and from the first client to identify queries and response portions, an interception mechanism that intercepts queries and response portions from the "other clients," an aggregation mechanism that aggregates those response portions, and a transmission mechanism that transmits at least a portion of the response to the first client.

(...continued)
mean "transparently" as that term is used in applicants' claims.

These elements of the claim cannot be read on Farber. The requesting computer in Farber must equate to the "first client." According to Farber, the requesting computer does not send its requests to multiple other clients as claimed. Each request is directed to a specific "origin server." That is the first additional failure of Farber relative to the claims.

And if the Farber repeater is the claimed acceleration server, then to meet the claims the repeater must intercept both queries from the first client to at least two other clients and responses to the first client from the at least two other clients. However, the repeater in Farber does not intercept queries from the first client to at least two other clients. Each repeater receives only queries from one client to a specific origin server, as reflected to it by a reflector. Nor can the reflector be the claimed acceleration server because all it does is instruct the first client to communicate with the repeater instead of with the origin server. This is a second additional failure of Farber relative to the claims.

It is telling that Farber, in describing that a repeater might seek missing data from other repeaters, describes that process as "peer[ing]" (Farber, column 3, line 48), but does not describe the operations relied on by the Examiner as peering.

Because the limitations of applicants' claims are missing from both Farber and Feigenbaum, the combination of Farber and Feigenbaum, even if proper, cannot show those limitations either.

For these reasons, applicants respectfully submit that independent claims 61, 75 and 80, and by extension all dependent claims, are patentable.

Conclusion

For the reasons set forth above, applicants respectfully submit that this application, as amended, is in condition for allowance. Reconsideration and prompt allowance of this application are respectfully requested.

Respectfully submitted,

/Jeffrey H. Ingerman/

Jeffrey H. Ingerman
Reg. No. 31,069
Attorney for Applicants
FISH & NEAVE IP GROUP
ROPES & GRAY LLP
Customer No. 1473
1251 Avenue of the Americas
New York, New York 10020-1105
Tel.: (212) 596-9000